

Abstract

A method of identifying a user-specified object contained in one or more images of a plurality of images that comprises the steps of defining regions of objects in the images, and computing a vector in respect of each of the regions based on the appearance of the respective region. The vector comprises a descriptor. The method further comprises vector quantizing the descriptors into clusters, storing the clusters as an index with the images in which they occur, defining regions of the user-specified object, computing a vector in respect of each of said regions based on the appearance of the regions, and vector quantizing the descriptors into one or more clusters. The index is searched and the clusters are compared with the contents of the index to identify which of the images contain the clusters so as to return the images containing the user-defined object.